

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: USING ERROR CHECKING BITS TO COMMUNICATE AN
ADDRESS OR OTHER BITS

APPLICANT: SRINIVAS KATAR AND LAWRENCE W. YONGE III

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EV 330508072 US

November 20, 2003
Date of Deposit

USING ERROR CHECKING BITS TO COMMUNICATE AN ADDRESS OR OTHER BITS

TECHNICAL FIELD

This invention relates to network communications.

BACKGROUND

Networks where multiple stations share a common medium require stations to coordinate
5 access to the medium. A variety of protocols like Carrier Sense Multiple Access (CSMA), Token
passing and Polling are used to achieve coordinate medium access. Stations in multiple access
networks exchange data using structured protocol entities called packets. These packets carry
addressing information required to identify the source and destination of the transmission.
Typically, each station in the network is identified by a unique address. These addresses are sent
10 in packets to uniquely identify the source and destination.

Certain network situations may require stations to broadcast these addresses. This can be
costly in some networks where broadcasting is inefficient and the addresses are large.

SUMMARY

In one aspect, the invention features a method of operating in a network in which a
15 plurality of stations communicate over a shared medium. The method comprises transmitting a
group of bits from a transmitting station to a receiving station, wherein the group of bits
comprises information bits and error checking bits, and communicating additional bits from the
transmitting station to the receiving station without transmitting the additional bits.
Communicating the additional bits comprises, at the transmitting station, deriving the error
20 checking bits from the additional bits using an error checking process; and at the receiving
station, using the same or a related error checking process to process candidate bits known to the
receiving station, and comparing the output of the error checking process to the error checking
bits received from the transmitting station to determine if the candidate bits correspond to the
additional bits.

25 Embodiments may include one or more of the following. The communicating of the
additional bits may further comprise, at the transmitting station, deriving the error checking bits
from the additional bits and the information bits using the error checking process, and at the

receiving station, using the same or a related error checking process to process the candidate bits and the information bits received from the transmitting station, comparing the output of the error checking process result to the error checking bits received from the first station to determine if the candidate bits correspond to the additional bits. The additional bits may comprise an address identifying one or both of the transmitting and receiving stations. The group of bits may constitute a packet, the information bits may comprise a source address and data, and the additional bits may comprise a destination address identifying the receiving station. The group of bits may constitute a packet, the information bits may comprise a source address and data, and the additional bits may comprise at least a portion of a destination address at least partially identifying the receiving station. The group of bits may constitute a Request To Send (RTS), the information bits may comprise a source address and data, with the source address identifying the transmitting station, and the additional bits may comprise a destination address identifying the receiving station. The group of bits may constitute a RTS, the information bits may comprise a source address and data, with the source address identifying the transmitting station, and the additional bits may comprise at least a portion of a destination address at least partially identifying the receiving station. The group of bits may constitute a Clear To Send (CTS) transmitted in response to an RTS, and the additional bits may comprise one or both of the following: (1) a destination address identifying the receiving station and (2) a source address identifying the transmitting station. The group of bits may constitute an acknowledgement and the additional bits may comprise one or both of the following: (1) a source address identifying the transmitting station and (2) a destination address identifying the receiving station. The group of bits may constitute an acknowledgement and the additional bits may comprise one or both of the following: (1) at least a portion of a source address at least partially identifying the transmitting station and (2) at least a portion of a destination address at least partially identifying the receiving station. The group of bits may constitute an acknowledgement transmitted in response to a receipt of a packet, and the additional bits may comprise a set of bits from the packet that is sufficiently unique as to associate the acknowledgement with the packet. The information bits may comprise bits indicating which portions of a previously transmitted packet need to be retransmitted. The error checking process may comprise a cyclic redundancy check, and the error checking bits may comprise frame check sequence (FCS) bits associated with the cyclic redundancy check.

and the error checking bits may comprise frame check sequence (FCS) bits associated with the cyclic redundancy check.

In another aspect, the invention features a method of operating in a network in which a plurality of stations communicate over a shared medium that includes transmitting a group of bits from a transmitting station to a receiving station, wherein the group of bits comprises information bits and a substantially unique number. The method also includes acknowledging receipt of the group of bits at the receiving station by transmitting an acknowledgement containing the substantially unique number or number or containing data from which the substantially unique number can be derived.

Embodiments may include one or more of the following. The group of bits may include a destination address identifying the receiving station, and the acknowledgement may include the same destination address. The transmitting station may receive the acknowledgement and may determine whether it contains the same destination address and the same substantially unique number as included in the transmitted group of bits. The group of bits may constitute an RTS and the acknowledgement may constitute a CTS. The substantially unique number may comprise a pseudorandom number.

These and other embodiments may have one or more of the following advantages. Packets may be transmitted and received by the intended destinations without requiring the transfer of complete source and destination addresses. This can be used to improve packet efficiency. For example, standard error checking techniques can be used to encode source and destination address information in packets without explicitly transmitting the source and destination address information. Various types of error checking techniques may be used, including cyclic redundancy code checking as well as checksum. For Request to Send (RTS) and Clear to Send (CTS), the communication overhead is reduced by not requiring transferring both source and destination addresses in the RTS message.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 is a schematic of a network configuration.

FIGS. 2-5 are data structure diagrams of transmitted and received data..

DETAILED DESCRIPTION

5 There are a great many possible implementations of the invention, too many to describe herein. Some possible implementations that are presently preferred are described below and in the Appendix. It cannot be emphasized too strongly, however, that these are descriptions of implementations of the invention, and not descriptions of the invention, which is not limited to the detailed implementations described in this section but is described in broader terms in the
10 claims.

 As shown in FIG. 1, network configuration 2 includes communications medium 3 and network 4 in which electronic devices 6, 8, and 10 (e.g., audiovisual equipment) communicate over medium 3. Electronic devices 6, 8, and 10 include media access controllers (MAC) 12, 14, and 16 that manage communication access to the network 4 for electronic devices 6, 8, and 10,
15 respectively. MACs 12, 14, and 16 implement the data link layer and connect to the physical layer of the Open Systems Interconnection (OSI) network architecture standard. In a general sense, MACs 12, 14, and 16 represent stations on network 4 that send messages to one another over medium 3. Communications medium 3 is a physical communication link between electronic devices 6, 8, and 10, and may include optical fiber, coaxial cable, unshielded twisted
20 pair, in addition to other media such as power lines. Electronic devices 6, 8, and 10 communicate with one another based on requirements of software applications running on electronic devices 6, 8, and 10. This communication creates traffic of messages on network 4.

 In some examples, medium 3 has a limited bandwidth due to limited physical capacity or electronic noise. Messages (e.g., groups of bits such as packets or ACKs) sent between MACs
25 include at least some information bits (e.g., packet body, payload, frame control information, addresses). Some messages between MACs 12, 14, and 16 also convey additional bits of information (e.g., a destination address) without sending the additional bits. The additional bits are conveyed by deriving error checking bits (e.g., CRC or checksum) from the additional bits at the transmitting station, transmitting the error checking bits with the information bits, and then
30 using an error checking process at a receiving station to derive error checking bits from

candidate bits (e.g. the receiving station's address). If the error checking bits derived at the receiving station compare favorably (e.g., match) to the received error checking bits, then the receiving station can safely conclude that the candidate bits are the same as the additional bits used by the transmitter (e.g., that the destination address is the receiver's address).

5 Communication error checking can be done in various ways. E.g., a transmitting station may use an error checking process to encode a block of data to be transmitted into a substantially unique code (which could be termed error checking bits) that is substantially unique to the communication and send the unique code along with the block of data to a receiving station. The receiving station then processes the received block of data and unique code using the same or
10 related error checking process to determine whether the received block of data has been received without error. Common error checking processes include cyclic redundancy code (CRC) checking and checksum.

 Cyclic redundancy code (CRC) checking is a method of checking for errors in data that have been transmitted on a communications link. For example, a transmitting station (e.g., MAC
15 12) in network configuration 2 typically applies a 8-, 16- or 32-bit polynomial to a block of data that is to be transmitted to a receiving station (e.g., MAC 14) and appends the resulting cyclic redundancy code to the block. The receiving station (e.g., MAC 14) applies the same polynomial to the data and each receiving station compares its result with the result appended by the transmitting station. If the result matches with the appended CRC, the receiving station can
20 assume that the block of data has been received successfully.

 Another typical error detection method for communication is the checksum method. A sum of data in a transmitted block (the checksum) is transmitted with the block of data so that the receiving station can compare a sum of the received block data with the received sum. If the sums match, then the receiving station can assume that the transmission of the block of data was
25 correctly received.

 MACs 12, 14, and 16 convey a block of data (information bits, e.g., a packet body or payload) as well as extra information (additional bits, e.g., a destination address) without using substantially more bandwidth than would have been necessary for conveying the block of data. This is accomplished by deriving error checking bits from the extra information and the block of
30 data but actually transmitting only the block of data and the error checking bits (alternatively, some portion, but not all, of the extra information might be transmitted). Although the extra

information is not transmitted, each station (e.g., MACs 14 and 16) receiving the message can perform an error checking process using candidate bits (e.g., the receiving station's address) to test whether the candidate bits correspond to the extra information (e.g., whether the receiving station's address matches the destination address). The receiving station may accomplish this by
5 appending the candidate bits (e.g., its address) to the received block of data message, and then applying the error checking process (the same one as used in the transmitter, or a related one that should produce a result from which it can be determined whether the candidate bits compare favorably to the additional bits). Various types of extra information (additional bits) can be conveyed in this manner, e.g., destination or source addresses, identification of a packet, or other
10 information from a finite universe of information known to the receiving station.

As shown in FIG 2, MAC 12 has a task of sending a message frame 102 with a block of data 104 to MAC 14 via medium 3 without directly specifying a destination address 108 for MAC 14. Both MAC 14 and MAC 16 will receive the message and MAC 14 should behave as if the message 102 is intended for MAC 14 while MAC 16 should function as if the message frame
15 102 is not intended for MAC 16. In the message frame 102, MAC 12 sends block of data 104 and error checking code 106. MAC 12 generates the error checking code 106 as if the block of data 104 were sent with the address 108 for MAC 14. Upon receipt of the message frame 102, MAC 14 determines that the message frame 102 is intended for MAC 14, and MAC 16 determines that the message is not intended for MAC 16. MAC 14 and MAC 16 make this
20 determination by inserting each receiving station's address 112 into an augmented message 110 (containing the received block of data and the receiving station's address) and using error checking code 106 against the augmented message 110 to check for errors. If there are no errors using this augmented message 110, then the inserted address 112 is assumed to be the destination address 108. In this example, MAC 14 has no errors and MAC 16 has errors.

Another example of matching some information in the message with an address of the receiving MAC occurs when MACs 12, 14, and 16 send Request to Send/Clear to Send (RTS/CTS) message pairs between each other. RTS/CTS is used when a MAC (e.g., 12) needs to know that another MAC (e.g., 14) is listening before sending data. As shown in FIG 3, MAC 12 sends a RTS message frame 202 to medium 3 to check if MAC 14 is listening before MAC 12
25 sends data to MAC 14. In this example, the RTS message frame 202 includes a data portion 204 and an error code 206 for error checking at the receiving station (e.g., MAC 14). In one

example, the error code 206 is a CRC resulting from applying a 8-, 16- or 32-bit polynomial to the data portion 204. The data portion 204 includes a source address 210 of MAC 12. MACs 14 and 16 receive the RTS message on medium 3. Both MAC 14 and 16 read the data portion 204 and the error code 206 of the message 202. In the example where error code 206 is a CRC, both
5 MAC 14 and 16 apply the 8-, 16- or 32-bit polynomial to an augmented message 212 representing the combination of the received data portion 204 and address 214 of either MAC 14 or 16. Both MAC 14 and 16 check for errors using the augmented message 212. MAC 14 finds no error in the augmented message 212 so MAC 14 detects that the RTS message is intended for MAC 14. By making this successful match, MAC 14 not only knows that the RTS message 202
10 is intended for MAC 14 but also that the message was received without transmission errors. MAC 14 responds to the RTS by sending a CTS message to medium 3 with MAC 12 as its intended recipient without including address 210 of MAC 12 in the message body. The CTS message has an error code that is derived from the data portion of the CTS message in addition to the destination address 210 of MAC 12. Alternatively, the CTS message can use one or both of
15 the source address (address of station generating the CTS) and destination address (address of station that is the intended recipient of the CTS) in the CRC computation. MACs 12 and 16 receive the CTS message on medium 3 and MAC 12 detects that the CTS message is intended for MAC 12 by augmenting the message with its address 210 and calculating that the augmented message is error free as described previously.

20 Other applications of this messaging technique include packet acknowledgement so that a transmitting station knows when its packets have been received or not. As shown in FIG. 4, MAC 12 sends a message frame 302 containing an acknowledgement of a receipt of a packet. MAC 14 originally sent the packet to MAC 12 so the packet acknowledgement conveys a source address identification for MAC 12. Upon receipt of this acknowledgement, MAC 14 can match
25 the acknowledgement to a packet sent by MAC 14 to MAC 12. The acknowledgement message 302 is intended for MAC 14 but both MAC 14 and 16 receive the message 302 on medium 3. MAC 12 generates the message 302 by generating error code 306 as if the source address 310 (address of MAC 12) and the destination address 308 (address of MAC 14) were part of a block of data for the message 302 without transmitting the source address 310 and the destination
30 address 308. Alternatively, the acknowledgement can also use only the destination address or only the source address in the CRC computation. The intended receiver knows the address of the

station that is supposed to send the response, and thus the acknowledgment mechanism can use one or both of the destination address and source address. Upon receipt of the message, both MAC 14 and MAC 16 insert their addresses into an augmented message 312 as potential destination address 316. Both MAC 14 and MAC 16 also insert addresses into an augmented message 312 of stations where MAC 14 and 16 previously sent unacknowledged packets as potential source addresses 314. MAC 14 and 16 then check their augmented messages 312 against the received error code 306 in the message to see if there is a match. In this case, MAC 14 has a match for the destination address 316 and a source address 314 for the packet previously sent to MAC 12.

FIG 5 shows one possible implementation of a communication method in which a substantially unique number (e.g., a pseudorandom number) is transmitted along with block of data, and a receiving station acknowledges receipt of the block of data by transmitting an acknowledgement containing the substantially unique number or containing data from which the substantially unique number can be derived. As described previously, MAC 12 wishes to determine if data can be sent to MAC 14. As shown in FIG 5, MAC 12 sends a RTS message 402 to MAC 14 with a substantially unique number 406, such as a pseudorandomly generated number, along with an address 408 of MAC 14 and other data 404. In response, MAC 14 sends back a CTS message 420 to MAC 12 with the same substantially unique number 406 and the address 408 of MAC 14 along with other data. Upon receipt of the CTS message 420, MAC 12 knows that it has received a CTS message corresponding to its RTS message 402 by a correct match of its original substantially unique number 406 and the received substantially unique number as well as the address 408 of MAC 12. Subsequently, MAC 12 can send data to MAC 14. In an alternative implementation a different number is returned, one that is derived by transforming number 406, but MAC 12 is aware of the transformation, and is able to ascertain that the returned number was derived from the number 406.

The normal cyclic redundancy code algorithm is typically as follows. First, a CRC is determined to be some number n bits long. To generate the CRC for a block of data, a transmitting station (e.g., MAC 12) left shifts the frame representing the block of data by n bits and then divides it by a predefined polynomial P . P is $n+1$ bits long. The remainder of the division is the FCS. MAC 12 appends the FCS to the frame representing the block of data. The result is the final frame forming the message body to transmit. On the receiving side, a receiving

station (e.g., MAC 14) receives the frame. MAC 14 then divides the frame by the predefined polynomial P . If the remainder of the division is zero then the transmission is error free. If the transmission is not zero then there is an error in the transmission.

Several different length CRCs are used in network configuration 2 by MACs 12, 14, and 16 to match destination addresses and check for errors using correct demodulation or decryption of data. Different lengths of these CRCs include CRC-8, CRC-16, and CRC-24.

The CRC-8 is calculated using the following standard generator polynomial of degree 8: $G(x) = x^8 + x^2 + x + 1$. The CRC-8 is the one's complement of the (modulo-2) sum of the following:

- The remainder of $x^k \times (x^7 + \dots + x^2 + x + 1)$ divided (modulo-2) by $G(x)$, where k is the number of bits in the calculation fields, and
- The remainder after multiplication of the contents (treated as a polynomial) of the calculation fields by x^8 and then division by $G(x)$.

The CRC-8 field is transmitted commencing with the coefficient of the highest-order term.

As a typical implementation, at the transmitter, the initial remainder of the division is preset to all ones and is then modified by division of the calculation fields by the generator polynomial $G(x)$. The one's complement of this remainder is transmitted, with the high-order bits first, as the CRC-8 field.

At the receiver, the initial remainder is preset to all ones and the serial incoming bits of the calculation fields and CRC-8, when divided by $G(x)$ in the absence of transmission errors, results in a unique nonzero remainder value. The unique remainder value is the polynomial:

$$x^7 + x^6 + x^5 + x^4 + x + 1$$

The CRC-16 is computed using the following standard generator polynomial of degree 16: $G(x) = x^{16} + x^{12} + x^5 + 1$. The CRC-16 is the one's complement of the (modulo-2) sum of the following:

- The remainder of $x^k \times (x^{15} + x^{14} + \dots + x^2 + x + 1)$ divided (modulo-2) by $G(x)$, where k is the number of bits in the calculation fields, and
- The remainder after multiplication of the contents (treated as a polynomial) of the calculation fields by x^{16} and then division by $G(x)$.

5 The CRC-16 field is transmitted commencing with the coefficient of the highest-order term.

As a typical implementation, at the transmitter, the initial remainder of the division is preset to all ones and is then modified by division of the calculation fields by the generator polynomial $G(x)$. The one's complement of this remainder is transmitted, with the high-order
10 bits first, as the CRC-16 field.

At the receiver, the initial remainder is preset to all ones and the serial incoming bits of the calculation fields and CRC-16, when divided by $G(x)$ in the absence of transmission errors, results in a unique nonzero remainder value. The unique remainder value is the polynomial: x^{12}

$$+ x^{11} + x^{10} + x^8 + x^3 + x^2 + x + 1$$

15 The CRC-32 is computed using the following standard generator polynomial of degree 32: $G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$. The CRC-32 is the one's complement of the (modulo-2) sum of the following:

- The remainder of $x^k \times (x^{31} + x^{30} + \dots + x^2 + x + 1)$ divided (modulo-2) by $G(x)$, where k is the number of bits in the calculation fields, and
- 20 ▪ The remainder after multiplication of the contents (treated as a polynomial) of the calculation fields by x^{32} and then division by $G(x)$.

The CRC-32 field is transmitted commencing with the coefficient of the highest-order term. As a typical implementation, at the transmitter, the initial remainder of the division is preset to all ones and is then modified by division of the calculation fields by the generator

polynomial $G(x)$. The one's complement of this remainder is transmitted, with the high-order bits first, as the CRC-16 field.

At the receiver, the initial remainder is preset to all ones and the serial incoming bits of the calculation fields and CRC-16, when divided by $G(x)$ in the absence of transmission errors, results in a unique nonzero remainder value. The unique remainder value is the polynomial: $x^{12} + x^{11} + x^{10} + x^8 + x^3 + x^2 + x + 1$.

Many other implementations of the invention other than those described above are within the invention, which is defined by the following claims.